

# Generalized Thompson sampling for sequential decision-making and causal inference

Pedro A. Ortega and Daniel A. Braun

March 20, 2013

## Abstract

Recently, it has been shown how sampling actions from the predictive distribution over the optimal action—sometimes called *Thompson sampling*—can be applied to solve sequential adaptive control problems, when the optimal policy is known for each possible environment. The predictive distribution can then be constructed by a Bayesian superposition of the optimal policies weighted by their posterior probability that is updated by Bayesian inference and causal calculus. Here we discuss three important features of this approach. First, we discuss in how far such Thompson sampling can be regarded as a natural consequence of the Bayesian modeling of policy uncertainty. Second, we show how Thompson sampling can be used to study interactions between multiple adaptive agents, thus, opening up an avenue of game-theoretic analysis. Third, we show how Thompson sampling can be applied to infer causal relationships when interacting with an environment in a sequential fashion. In summary, our results suggest that Thompson sampling might not merely be a useful heuristic, but a principled method to address problems of adaptive sequential decision-making and causal inference.

## 1 Introduction

In a research paper from 1933, Thompson studied the problem of finding out which one of two drugs was better when testing them on a patient population under the constraint that as few people as possible should be subjected to the inferior drug during the course of testing [1]. Given a current probability estimate  $P$  of one treatment being better than the other, he suggested that it might be a good idea to adjust the proportions of future test subjects that take the two drugs to the respective probabilities  $P$  and  $1 - P$ . This way one would not run into the danger of *permanently* cutting off all future test subjects from a potentially superior treatment that so far seems inferior due to statistical fluctuations, while only *temporarily* risking exposure to a potentially inferior drug for a smaller proportion of the population. Randomizing actions based on the probability that this action is believed to be optimal when faced with an unknown environment is now sometimes called *Thompson sampling*.

Thompson sampling is a form of *probability matching*. Probability matching has been extensively studied in both humans and animals when they make predictions in stochastic environments [2, 3]. Rather than consistently predicting the most likely outcome, experimental subjects tend to randomize their predictions based on the probabilities with which the respective events occur. When knowing the probabilities, this is clearly a suboptimal strategy. However, in the case of Thompson sampling it is important to note that the probabilities are not known. Nevertheless, one might argue that Thompson sampling is a suboptimal strategy, as Thompson’s problem can be thought of as a *bandit problem* [4], which is solved optimally by Gittin’s indices in the case of known prior probabilities and discounted rewards [5]. Most studies have therefore examined Thompson sampling as a heuristic in the context of bandit problems [6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16].

Recently, it was shown, however, that Thompson sampling can also be applied to solve a more general class of sequential adaptive control problems, provided that both an optimal policy and a predictive model is known for each possible environment [17]. When an environment is drawn randomly from the set of possible environments, the optimal policy can then be inferred on the fly by an adaptation process that is driven by actions sampled from the predictive distribution over the optimal policies. Here we study three characteristic features of such generalized Thompson sampling. First, we discuss in how far Thompson sampling can be regarded as a natural consequence of a Bayesian treatment of policy uncertainty. Second, we study convergence behavior when two adaptive Thompson sampling agents are coupled in a sequential fashion. Third, we show how this approach can be extended naturally to address the problem of causal induction when interacting with an unknown environment.

The paper is structured as follows. In Section 2 we clarify the problem statement and recapitulate the main result of [17]. In Section 3 we analyze the uncertainty faced by agents that are unable to compute the single best policy. In Section 4 we study interactions that arise when coupling two adaptive agents that employ Thompson sampling to determine their actions. In Section 5 we investigate how this approach can be applied to adaptive agents that need to discover the causal structure of their environment. Finally, we discuss in what sense Thompson sampling might provide a principled solution to adaptive decision-making problems.

## 2 Problem Statement

### 2.1 Preliminaries

We restrict the exposition to the case of discrete time with discrete stochastic observations and actions. Let  $\mathcal{O}$  and  $\mathcal{A}$  be two finite sets, the first being the *set of observations* and the second being the *set of actions*. We use  $a_{\leq t} := a_1 a_2 \dots a_t$  and  $a_{< t} := a_1 a_2 \dots a_{t-1}$  to simplify the notation of strings. We define the *set of interactions* as  $\mathcal{Z} := \mathcal{A} \times \mathcal{O}$ . The set of interaction strings of length  $t \geq 0$

is denoted by  $\mathcal{Z}^t$ . The set of all finite interaction strings is  $\mathcal{Z}^* := \bigcup_{t \geq 0} \mathcal{Z}^t$ , the set of infinite strings is  $\mathcal{Z}^\infty := \{w : w = a_1 o_1 a_2 o_2 \dots\}$ . The interaction string of length 0 is denoted by  $\epsilon$ .

Agents and environments are formalized as I/O systems. An *I/O system*  $\mathbf{Pr}$  is a probability distribution over interaction sequences  $\mathcal{Z}^\infty$ .  $\mathbf{Pr}$  is uniquely determined by the conditional probabilities

$$\mathbf{Pr}(a_t | a_{<t}, o_{<t}), \quad \mathbf{Pr}(o_t | a_{\leq t}, o_{<t})$$

for each  $a_1 o_1 \dots a_{t-1} o_{t-1} a_t \in \mathcal{Z}^*$ . An *interaction system*  $(P, Q)$  is a coupling of two I/O systems, where  $P$  is an agent and  $Q$  is an environment. Because the agent and the environment mutually influence each other, their actions and observations are conditioned by the previous interactions. Accordingly, the probability of an interaction string  $a_1 o_1 \dots a_T o_T$  is given by

$$\prod_{t=1}^T P(a_t | a_{<t}, o_{<t}) Q(o_t | a_{\leq t}, o_{<t}). \quad (1)$$

From the point of view of the agent  $P$ , the distribution  $P(a_t | a_{<t}, o_{<t})$  is a *policy* and captures the probability of producing action  $a_t \in \mathcal{A}$  given history  $a_1 o_1 \dots a_{t-1} o_{t-1}$ . The distribution  $P(o_t | a_{\leq t}, o_{<t})$  is the agent's *predictive model* of the environment, as it predicts the probability of the observation  $o_t \in \mathcal{O}$  given history  $a_1 o_1 \dots a_{t-1} o_{t-1} a_t$ . For the agent  $P$ , the sequence  $o_1 o_2 \dots$  provides its input stream and the sequence  $a_1 a_2 \dots$  is its output stream. In the case of the environment  $Q$  the roles are reversed, that is the sequence  $o_1 o_2 \dots$  is its output stream and the sequence  $a_1 a_2 \dots$  provides its input stream. The quintessential goal is to choose the agent's policy such that the resulting distribution over the interaction sequences (1) is *desirable*.

## 2.2 Policy Construction: Known Environment

If  $Q$  is known, then  $P$  can be equipped with a model that can perfectly predict its environment, that is  $P(o_t | a_{\leq t}, o_{<t}) = Q(o_t | a_{\leq t}, o_{<t})$  for all  $a_1 o_1 \dots a_t \in \mathcal{Z}^*$ . Moreover, a custom-made policy can be designed for  $P$  that produces desirable interaction sequences. Desirability is typically formalized by the economic theory of subjective expected utility (SEU) [18, 19], which stipulates that a decision maker's preferences over lotteries *can be thought of* as maximizing a SEU of the outcome. In the policy construction setting, this translates into the designer having a real-valued *utility function* giving rise to utilities  $U(a_{1:T}, o_{1:T})$  for each realization, and the *predictive model*  $P(o_t | a_{\leq t}, o_{<t})$ . The utility function quantifies the subjective desirability of a particular interaction string and the probabilities represent the subjective model of the environment. The *maximum expected utility principle* then states that the designer has to choose the policy such that it maximizes the expected utility

$$\mathbf{E}_P[U] := \sum_{a_{\leq T}, o_{\leq T}} P(a_{\leq T}, o_{\leq T}) U(a_{\leq T}, o_{\leq T})$$

where the probabilities  $P(a_{\leq T}, o_{\leq T})$  are the policy-prediction products

$$\begin{aligned} & \prod_{t=1}^T P(a_t | a_{<t}, o_{<t}) P(o_t | a_{\leq t}, o_{<t}) \\ &= \underbrace{\left\{ \prod_{t=1}^T P(a_t | a_{<t}, o_{<t}) \right\}}_{\text{policy}} \underbrace{\left\{ \prod_{t=1}^T P(o_t | a_{\leq t}, o_{<t}) \right\}}_{\text{prediction}}. \end{aligned} \quad (2)$$

The optimal policy is often computed by restating the problem recursively and then using dynamic programming to solve the Belmann optimality equations [20]. The policy and the prediction model are both subjective in the sense that they are unilaterally chosen by the designer. A policy choice explainable by this scheme is defined to be a *rational* choice. Choices that do not strictly obey the maximum SEU principle are *irrational*, or at best *bounded rational* [21].

### 2.3 Policy Construction: Unknown Environment

In general, the prediction model will not be equal to the generative law of the environment, that is,

$$P(o_t | a_{\leq t}, o_{<t}) \neq Q(o_t | a_{\leq t}, o_{<t}).$$

and consequently the true expected utility is in general not equal to the SEU:

$$\mathbf{E}_{\text{Pr}}[U] \neq \mathbf{E}_P[U]. \quad (3)$$

One of the most interesting cases where the prediction model and the true generative law of the environment do not match is when the designer is uncertain about the latter. Formally, the designer expresses his uncertainty by introducing a random variable  $\theta$  that indexes the class of potential environments  $Q_\theta$ . More specifically, he has a class of prediction models and policies

$$B(o_t | \theta, a_{\leq t}, o_{<t}), \quad B(a_t | \theta, a_{<t}, o_{<t}), \quad (4)$$

such that for every possible environment indexed by  $\theta$  there is a perfectly fitting predictor  $B(o_t | \theta, a_{\leq t}, o_{<t}) = Q_\theta(o_t | a_{\leq t}, o_{<t})$  and a desirable custom-built policy  $B(a_t | \theta, a_{<t}, o_{<t})$ . Moreover, the designer believes that  $Q_\theta$  is drawn with probability  $B(\theta)$  from a set  $\Theta$  of possible environments before the interaction starts, where  $\Theta$  is assumed to be discrete for simplicity.

#### 2.3.1 Decision-theoretic Problem Formulation

In order to stay within the framework of subjective expected utility one has to reduce the problem of the unknown environment to a problem with known environment. Such a “known” environment can be created from a set of possible environments as a new “super-environment” by marginalizing over the

parameter of the possible environments, thus, obtaining the *Bayesian mixture distribution* [22]

$$\begin{aligned} B(o_{\leq T}|a_{\leq T}) &= \sum_{\theta} B(\theta) B(o_{\leq T}|\theta, a_{\leq T}) \\ &= \sum_{\theta} B(\theta) \prod_{t=1}^T B(o_t|\theta, a_{\leq t}, o_{< t}). \end{aligned}$$

The adaptive control problem is then solved by equating the prediction model  $P$  with the Bayesian predictive distribution over the observations,

$$P(o_t|a_{\leq t}, o_{< t}) := B(o_t|a_{\leq t}, o_{< t}),$$

and then choosing the policy that maximizes the SEU as in the case of a known environment. This procedure effectively enlarges the space of possible environments to the convex hull  $H(\Theta)$  spanned by the prediction models in  $\Theta$ , i.e. an “environment” is any convex combination of distributions indexed by  $\theta$ . Each  $\theta \in H(\Theta)$  thus corresponds to a Bayesian mixture over the environments in  $\Theta$ , but with a different prior. If the true generative law coincides with one of the models, i.e. there is a  $\theta^* \in \Theta$  such that

$$Q(o_t|a_{\leq t}, o_{< t}) = B(o_t|\theta^*, a_{\leq t}, o_{< t}), \quad (5)$$

then the predictive model will converge to the true generative law with Probability one, that is

$$P(o_t|a_{\leq t}, o_{< t}) \rightarrow Q(o_t|\theta^*, a_{\leq t}, o_{< t})$$

as  $t \rightarrow \infty$ . Since the policy choice is optimal, this scheme directly bypasses the *exploration-exploitation dilemma* [23]. The important point, however, is that in order to express the uncertainty over the environment’s probability law, the designer had to introduce a belief model that *compiles* into an actual prediction model. Thus, the policy construction in the SEU problem statement for unknown environments proceeds in two steps: first, a Bayesian mixture environment is created, and second a utility-maximizing policy is found on this mixture environment.

### 2.3.2 Probabilistic Problem Formulation

An alternative formulation of the problem statement for unknown environments is a one-step procedure that essentially stays within (Bayesian) probability calculus. In this case actions are treated as random variables and our goal is to determine a distribution  $P(a_t|\hat{a}_{< t}, o_{< t})$  that tells us how to act depending on past actions  $\hat{a}_{< t}$  and past observations  $o_{< t}$ . The distribution  $P(a_t|\hat{a}_{< t}, o_{< t})$  is the pendant to the observational distribution  $P(o_t|\hat{a}_{\leq t}, o_{< t})$  used for prediction. The only caveat is that past probabilistic actions, unlike past observations, have to be marked as causal interventions—denoted by  $\hat{o}$  in causal calculus [24].

Given the models (4), both distributions can then be expressed as mixture distributions

$$\begin{aligned} P(o_t|\hat{a}_{\leq t}, o_{< t}) &= \sum_{\theta} B(o_t|\theta, a_{\leq t}, o_{< t}) B(\theta|\hat{a}_{\leq t}, o_{< t}) \\ P(a_t|\hat{a}_{< t}, o_{< t}) &= \sum_{\theta} B(a_t|\theta, a_{< t}, o_{< t}) B(\theta|\hat{a}_{< t}, o_{< t}), \end{aligned}$$

where the posterior is given by

$$B(\theta|\hat{a}_{< t}, o_{< t}) = \frac{B(o_{t-1}|\theta, a_{< t-1}, o_{< t-1}) B(\theta|\hat{a}_{< t-1}, o_{< t-1})}{\sum_{\theta'} B(o_{t-1}|\theta', a_{< t-1}, o_{< t-1}) B(\theta'|\hat{a}_{< t-1}, o_{< t-1})}.$$

Sampling from  $P(a_t|\hat{a}_{< t}, o_{< t})$  is equivalent to sampling a random belief  $\theta$  from the posterior  $B(\theta|\hat{a}_{< t}, o_{< t})$  and then acting according to  $B(a_t|\theta, a_{< t}, o_{< t})$ . This corresponds to a generalized Thompson sampling procedure, where first a random belief is sampled and then the optimal policy with respect to this belief is executed. Effectively, the posterior is also the only place where causal calculus plays a role, as can be seen by the absence of the likelihood  $B(\hat{a}_{t-1}|\theta, a_{< t-1}, o_{< t-1})$ , which is equal to one. Intuitively, the reason for this is that the agent can be surprised about his past observations and learn from them, but he cannot be surprised about his own actions chosen by himself in the past. Past actions do not provide any information about the environment. As will be explained in more detail below, causal calculus deals exactly with inference problems when some random variables are intervened or set by the inference machine itself. Importantly, this result is obtained solely by applying basic probability and causal calculus.

### 3 Policy Uncertainty

While subjective expected utility is formally appealing as a principle for the construction of adaptive agents, its strict application is in practice often problematic. This is mainly due to two reasons:

1. *Computational complexity.* The computations required to find the optimal solution (for instance, the computational complexity of solving the Bellman optimality equations) are prohibitive in general and scale exponentially with the length of the horizon: the time complexity of the search algorithm is  $\mathcal{O}(|\mathcal{A}| \times |\mathcal{O}|^T)$ . The problem is tractable only in very special cases under assumptions that reduce the effective size of the problem.
2. *Causal precedence of policy choice.* The choice of the policy has to be made before the interaction with the environment starts. That is, an agent has to have a unique optimal policy before it has even interacted once with the environment. An optimal policy constructed by the maximum SEU principle is therefore a very risky bet, as a lot of resources have to be spent before any evidence exists that the underlying model or prior is adequate.

Because of these two reasons, it is practically often impossible or questionable to apply the maximum SEU principle. In the following, we investigate how to weaken the formal assumptions of the policy construction method.

### 3.1 Policy search

Given a problem specification in terms of the predictive model and the utility function, the task of a policy search method is to calculate a policy that approximates the optimal policy.

More specifically, let  $\pi$  be a parameter in a set  $\Pi$  indexing the set of candidate policies

$$B(a_t|\pi, a_{<t}, o_{<t}) \quad (6)$$

analogous to the prediction models (4) indexed by  $\theta \in \Theta$ . *Then, in the most general case, a policy search method returns a probability distribution  $B(\pi)$  over  $\Pi$  representing the uncertainty over the optimal policy parameters.* If the algorithm solves the maximum SEU problem, then the support of this distribution will exclusively cover the set of optimal policies  $\Pi^* \subset \Pi$ . Otherwise there remains uncertainty over the optimal policy parameters.

Policies can also be parameterized in terms of the predictive model. In particular, we will assume that for each  $\theta \in \Theta$  there is a *known* optimal policy  $\pi \in \Pi$ , such that one can construct a function  $b : \Theta \rightarrow \Pi$  that maps each  $\theta$  into some  $\pi \in \Pi$ . Uncertainty over the environment can then directly be translated into policy uncertainty, such that any point in the convex hull  $H(\Theta)$  can be mapped to a corresponding point in the convex hull  $H(\Pi)$  spanned by the policies  $\pi \in \Pi$ .

### 3.2 The Exploration-Exploitation Trade-Off

Many policy search methods do not explicitly deal with the uncertainty over the policy parameters. Some methods only return a point estimate  $\hat{\pi} \in \Pi$ . It is obvious that the greedy usage of the estimate  $\hat{\pi}$  leads to sub-optimal performance, since for all  $\hat{\pi}$  that are not in the set  $\Pi^*$  of optimal policies, one has that

$$\mathbf{E}_B[U|\pi^*] > \mathbf{E}_B[U|\hat{\pi}]$$

where

$$\mathbf{E}_B[U|\pi] = \sum_{a \leq T, o \leq T} B(a_{\leq T}, o_{\leq T}|\pi) U(a_{\leq T}, o_{\leq T})$$

is the SEU with respect to the policy parameter  $\pi \in \Pi$ . For instance, reinforcement learning algorithms [4] start from a randomly initialized point estimate  $\hat{\pi}_0$  of the optimal policy and then generate refined point estimates  $\hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3, \dots$  in each time step  $t = 1, 2, 3, \dots$  using the data provided by experience. In order to converge to the optimal policy, these algorithms have to deal with the exploration-exploitation trade-off. This means that the agents cannot just greedily act according to these point estimates; instead, they have to produce

explorative actions as well, that is, actions that deviate from the current estimate of the optimal policy—for instance producing optimistic actions based on UCB [25, 26].

Let  $B_t(\pi)$  denote the posterior distribution over the optimal policy at time  $t$ . Then,

$$B_t(\pi) = B_t(b(\theta)) = B_t(b^{-1}(\pi)),$$

where  $b^{-1}(\pi) \subset \Theta$  is the pre-image of  $\pi$ . Hence, this shows that finding the optimal policy amounts to finding the pre-image of  $\pi^*$ , such that the distribution over the policy space becomes the delta function

$$B_t(\pi) = \begin{cases} 1 & \text{if } \theta^* \in b^{-1}(\pi), \\ 0 & \text{else} \end{cases}$$

where  $\theta^* \in \Theta$  is the true prediction model defined in (5). This highlights the essence of the exploration-exploitation trade-off: any action issued by the agent has to respect the uncertainty over the policy parameter—otherwise they are biased. In particular, if the agent acts greedily (i.e. it treats the estimate  $\hat{\pi}$  as if it were the true policy parameter) then it is overfitting the experience; likewise, an agent having excessive uncertainty is underfitting. From a frequentist point of view, this reveals that the exploration-exploitation trade-off is nothing else but the *bias-variance trade-off* [27] in policy space. This suggests that just like Bayesian modeling naturally balances the bias-variance trade-off by creating estimators that are probability distributions instead of point estimates, Bayesian modeling of the exploration-exploitation trade-off leads to a Bayes-causal solution for generalized Thompson sampling.

### 3.3 Bayes-Causal Solution

It is important to note that the concept of the bias is conditional on the true parameter—which is unknown when the designer is uncertain about the environment. This is not a problem from a Bayesian point of view, because the best estimator of the policy parameter is its posterior distribution [28]. Hence, instead of dealing with the exploration-exploitation dilemma by introducing explorative actions, one can directly use the posterior distribution over the policy parameter as an estimate.

To see how to do this, note that, by virtue of the mapping  $\pi = b(\theta)$ , the policies are independent of the policy parameter when the environment parameter is known:

$$B(a_t|b(\theta), a_{<t}, o_{<t}) = B(a_t|\theta, a_{<t}, o_{<t}).$$

Hence, each  $\theta \in \Theta$  indexes a dynamical model given by the distributions over interaction sequences

$$B(a_{\leq T}, o_{\leq T}|\theta) = \prod_{t=1}^T B(a_t|\theta, a_{<t}, o_{<t})B(o_t|\theta, a_{\leq t}, o_{<t}).$$



Given the dynamical models and their prior probabilities, the designer can form the Bayesian mixture model

$$B(a_{\leq T}, o_{\leq T}) = \sum_{\theta} B(\theta) B(a_{\leq T}, o_{\leq T} | \theta) \quad (7)$$

where the sum spans all the parameters in  $\Theta$ . The mixture models in  $H(\Theta) \setminus \Theta$  need not be considered here, since it is assumed that  $\theta^* \in \Theta$ .

### **Actions as Causal Interventions**

The designer can directly use the probabilistic model (7) to characterize an agent with policy uncertainty. There is a *caveat* though when actions are treated as random variables. It is clear that the observations produced by the environment update the agent’s state of knowledge about the environment. However, the actions are set by the agent itself and hence they do not provide information about the environment.

The theory that deals with the distinction between exogenous and endogenous information is *statistical causality* [29, 30]. Observations change the information state by regular Bayesian conditioning, whereas actions constitute causal interventions followed by Bayesian conditioning. To calculate the effect of an intervention, the causal model, i.e. the unique factorization of the joint distribution into conditional probabilities matching the causal dependencies over the random variables, is required to be known. In our setup, this is straightforward: first, the environment secretly chooses a true parameter  $\theta^* \in \Theta$ , and then the interactions  $a_1, o_1, a_2, o_2, \dots$  follow chronologically.

Formally, this means that the posterior probabilities over the environment parameters are given by

$$B(\theta | \hat{a}_{<t}, o_{<t})$$

rather than the more familiar expression  $B(\theta | a_{<t}, o_{<t})$ , where the “hat”-notation  $\hat{a}_t$  denotes a causal intervention [31].

For our needs, it is enough to consider the following simple method to calculate the effect of causal interventions:

1. Expand the probabilities in terms of the joint distribution.
2. Rewrite the joint distribution as the causal factorization.
3. Remove the intervention tags from the intervened random variables that are in the probability conditions.
4. Replace each conditional probability having an intervened variable in the argument by a delta function over its chosen value.

Applying these four steps to the posterior probabilities over the environment

parameters yields

$$\begin{aligned}
& B(\theta|\hat{a}_{<t}, o_{<t}) \\
& \stackrel{(1.)}{=} \frac{B(\theta, \hat{a}_{<t}, o_{<t})}{\sum_{\theta'} B(\theta', \hat{a}_{<t}, o_{<t})} \\
& \stackrel{(2.)}{=} \frac{B(\theta) \prod_{k=1}^t B(\hat{a}_k|\theta, \hat{a}_{<k}, o_{<k}) B(o_k|\theta, \hat{a}_{\leq k}, o_{<k})}{\sum_{\theta'} B(\theta') \prod_{k=1}^t B(\hat{a}_k|\theta', \hat{a}_{<k}, o_{<k}) B(o_k|\theta', \hat{a}_{\leq k}, o_{<k})} \\
& \stackrel{(3.)}{=} \frac{B(\theta) \prod_{k=1}^t B(\hat{a}_k|\theta, a_{<k}, o_{<k}) B(o_k|\theta, a_{\leq k}, o_{<k})}{\sum_{\theta'} B(\theta') \prod_{k=1}^t B(\hat{a}_k|\theta', a_{<k}, o_{<k}) B(o_k|\theta', a_{\leq k}, o_{<k})} \\
& \stackrel{(4.)}{=} \frac{B(\theta) \prod_{k=1}^t B(o_k|\theta, a_{\leq k}, o_{<k})}{\sum_{\theta'} B(\theta') \prod_{k=1}^t B(o_k|\theta', a_{\leq k}, o_{<k})}. \tag{8}
\end{aligned}$$

This equation shows that beliefs are updated only using the observations, and that actions are treated “as if they were known beforehand”, thus providing no evidence.

Likewise, note that

$$B(a_t|\theta, \hat{a}_{<t}, o_{<t}) = B(a_t|\theta, a_{<t}, o_{<t}). \tag{9}$$

Using (8) and (9), we get the probability of issuing action  $a_t \in \mathcal{A}$ :

$$B(a_t|\hat{a}_{<t}, o_{<t}) = \sum_{\theta} B(a_t|\theta, a_{<t}, o_{<t}) B(\theta|\hat{a}_{<t}, o_{<t}). \tag{10}$$

The important fact about (10) is it was derived only from probability theory and causal calculus by assuming policy uncertainty over a set of policies. We can therefore define the policy and prediction models of the agent as

$$\begin{aligned}
P(a_t|a_{<t}, o_{<t}) &:= B(a_t|\hat{a}_{<t}, o_{<t}) \\
P(o_t|a_{\leq t}, o_{<t}) &:= B(o_t|\hat{a}_{\leq t}, o_{<t}). \tag{11}
\end{aligned}$$

The construction of an adaptive agent with policy uncertainty then proceeds analogous to a Bayesian inference process.

1. First, we define a set of prediction models  $B(o_t|\theta, a_{<t}, o_{<t})$  and policy models  $B(a_t|\theta, a_{\leq t}, o_{<t})$ , where each policy model is optimal for a particular environment  $\theta$ . In the case of inference the latent random variable  $\theta \in \Theta$  corresponds to the hypothesis.
2. Second, we choose some prior probabilities  $B(\theta)$  to model our prior uncertainty.
3. Third, we use the distribution  $B(a_t|\hat{a}_{<t}, o_{<t})$  as the agent’s adaptive policy  $P(a_t|a_{<t}, o_{<t})$ , and the distribution  $B(o_t|\hat{a}_{\leq t}, o_{<t})$  as the agent’s adaptive predictor  $P(o_t|a_{\leq t}, o_{<t})$ .

Thus, Thompson sampling is used in every time step to sample an action  $a_t$  from the predictive distribution  $B(a_t|\hat{a}_{<t}, o_{<t})$ .

## 4 Convergence & Co-Adaptation

In [17], the limit behavior of a Thompson sampling agent (11) was investigated. Assuming that there exists a belief  $P(o_t|\theta, a_{\leq t}, o_{< t})$  that perfectly models the environment  $Q_\theta$  such that  $P(o_t|\theta, a_{\leq t}, o_{< t}) = Q_\theta(o_t|a_{\leq t}, o_{< t})$ , then the agent (11) converges in the sense that  $P(a_t|a_{< t}, o_{< t}) \rightarrow P(a_t|\theta, a_{< t}, o_{< t})$  almost surely as  $t \rightarrow \infty$  if the interaction system  $(P, Q)$  fulfills certain ergodicity requirements and all policies  $P(a_t|\theta, a_{< t}, o_{< t})$  are consistent. Roughly speaking, the first requirement ensures that the agent can recover from any initial mistakes, and the second requirement ensures that all predictors  $B(o_t|\theta, a_{\leq t}, o_{< t})$  that make the same predictions for the tail of the observation sequence are coupled to the same policy  $B(a_t|\theta, a_{< t}, o_{< t})$ . Thus, the same beliefs imply the same behaviors.

But what happens if the environment is also adaptive? As long as the agent has a model  $B(o_t|\theta, a_{\leq t}, o_{< t})$  that captures the adaptive behavior of the environment nothing fundamentally changes. However, the agent might not have a model about the *adaptive* behavior of the environment, while still having a pretty good idea about the environment's preferences. This is typically the case in game theory [32, 33], where the agent knows the other agent's best response function, but has no model of the other agent's adaptive behavior. In the simplest one-shot simultaneous move games the best response functions are given by

$$\begin{aligned} BR_1[Q(o)] &= \arg \max_{P(a)} \sum_{a,o} P(a)Q(o)U(a,o) \\ BR_2[P(a)] &= \arg \max_{Q(o)} \sum_{a,o} P(a)Q(o)V(a,o), \end{aligned}$$

where player 1's best response  $BR_1$  is a distribution  $P(a)$  over actions  $a$  that depends on agent 2's probability  $Q(o)$  of emitting  $o$ , and where  $U$  and  $V$  are the payoff functions for player 1 and 2 respectively. A Nash equilibrium  $(P^*(a), Q^*(o))$  is a fix point of these coupled equations [34], where each individual player has no incentive to change his distribution, that is

$$\begin{aligned} \sum_{a,o} P^*(a)Q^*(o)U(a,o) &\geq \sum_{a,o} P(a)Q^*(o)U(a,o) \quad \forall P(a) \\ \sum_{a,o} P^*(a)Q^*(o)V(a,o) &\geq \sum_{a,o} P^*(a)Q(o)V(a,o) \quad \forall Q(o). \end{aligned}$$

How such equilibria are reached is not subject in classic equilibrium game theory. But, if such games are repeated over and over again, evolutionary game theory suggests that these equilibria appear as fix points of adaptation dynamics like the replicator equations [35]. As Bayesian inference can also be viewed as some kind of replicator dynamics [36, 37], this provides an interesting starting point to study the emergence of Nash equilibria when two adaptive Thompson sampling agents interact.

When both agents are adaptive according to (11), we can decompose the

policies  $P(a)$  and  $Q(o)$  into mixture distributions, such that

$$\begin{aligned} P(a) &= \sum_{\theta} P(a|\theta)P(\theta) \\ Q(o) &= \sum_{\xi} Q(o|\xi)Q(\xi), \end{aligned}$$

where  $P(a|\theta) = BR_1[Q(o)]$  and  $Q(o|\xi) = BR_2[P(a)]$  and  $P(\theta)$  and  $Q(\xi)$  are the prior distributions over the possible behaviors. Moreover, both agents will have predictive models over the other agent's behavior, such that

$$\begin{aligned} P(o) &= \sum_{\theta} P(o|\theta)P(\theta) \\ Q(a) &= \sum_{\xi} Q(a|\xi)Q(\xi), \end{aligned}$$

where  $P(o|\theta) = Q(o|\xi)$  and  $Q(a|\xi) = P(a|\theta)$ . In the following we will assume that there exists at least one pair  $(\theta^*, \xi^*)$  where the model  $P(o|\theta^*)$  perfectly matches  $Q(o|\xi^*)$  such that  $P(o|\theta^*) = Q(o|\xi^*)$  and at the same time the model  $Q(a|\xi^*)$  perfectly matches  $P(a|\theta^*)$  such that  $P(a|\theta^*) = Q(a|\xi^*)$ . In this case both agents can predict the other agent's behavior, which means that there will be no drive to change the posteriors  $P(\theta|D)$  and  $Q(\xi|D)$ , given some past experience  $D$ . Then, both agents should “lock in” when their posteriors are sufficiently close to  $\delta_{\theta, \theta^*}$  and  $\delta_{\xi, \xi^*}$ . Formally, we define a pair  $(\xi^*, \theta^*)$  to be a strict Nash equilibrium if

$$\begin{aligned} \sum_o Q(o|\xi^*) \log \frac{Q(o|\xi^*)}{P(o|\theta^*)} &< \sum_o Q(o|\xi^*) \log \frac{Q(o|\xi^*)}{P(o)} \quad \forall P(o) \\ \sum_a P(a|\theta^*) \log \frac{P(a|\theta^*)}{Q(a|\xi^*)} &< \sum_a P(a|\theta^*) \log \frac{P(a|\theta^*)}{Q(a)} \quad \forall Q(a). \end{aligned}$$

This corresponds to a lock-in of the predictive models that drive the adaptation process.

In order to study convergence to  $(\theta^*, \xi^*)$ , we determine the difference in relative entropy between the predictive distribution  $P(o)$  at time  $t$  and the generative distribution  $Q(o|\xi^*)$  and the relative entropy between the predictive distribution  $P(o|D)$  at time  $t+1$  and the generative distribution  $Q(o|\xi^*)$ —after observing  $D$  at time step  $t$ . This difference can serve as a Lyapunov function to show convergence, if we require that

$$\Delta KL = \sum_D Q(D|\xi^*) D_{KL}(Q(o|\xi^*) \| P(o|D)) - D_{KL}(Q(o|\xi^*) \| P(o)) < 0.$$

Given the predictive distribution  $P(o) = \sum_{\theta} P(o|\theta)P(\theta)$  at time  $t$  and the

predictive distribution  $P(o|D) = \sum_{\theta} P(o|\theta)P(\theta|D)$  at time  $t + 1$ , we get

$$\begin{aligned}
\Delta KL &= \sum_D Q(D|\xi^*) \sum_o Q(o|\xi^*) \left[ \log \frac{Q(o|\xi^*)}{\sum_{\theta} P(o|\theta)P(\theta|D)} - \log \frac{Q(o|\xi^*)}{\sum_{\theta} P(o|\theta)P(\theta)} \right] \\
&= \sum_D Q(D|\xi^*) \sum_o Q(o|\xi^*) \log \frac{\sum_{\theta} P(o|\theta)P(\theta)}{\sum_{\theta} P(o|\theta)P(\theta|D)} \\
&\leq \sum_D Q(D|\xi^*) \sum_o Q(o|\xi^*) \left[ \log \frac{\sum_{\theta} P(o|\theta)P(\theta)}{P(o|\theta^*)P(\theta^*|D)} \right] \\
&= \sum_o Q(o|\xi^*) \log \frac{\sum_{\theta} P(o|\theta)P(\theta)}{P(o|\theta^*)} - \sum_D Q(D|\xi^*) \log P(\theta^*|D) \\
&= \sum_o Q(o|\xi^*) \log \frac{\sum_{\theta} P(o|\theta)P(\theta)}{P(o|\theta^*)} - \sum_D Q(D|\xi^*) \log \frac{P(D|\theta^*)P(\theta^*)}{\sum_{\theta'} P(D|\theta')P(\theta')} \\
&= \underbrace{\sum_o Q(o|\xi^*) \log \frac{\sum_{\theta} P(o|\theta)P(\theta)}{P(o|\theta^*)}}_{<0 \text{ (Nash property)}} + \underbrace{\sum_D Q(D|\xi^*) \log \frac{\sum_{\theta'} P(D|\theta')P(\theta')}{P(D|\theta^*)}}_{<0 \text{ (Nash property)}} - \underbrace{\log P(\theta^*)}_{\geq 0}
\end{aligned}$$

If the prior weight  $P(\theta^*)$  is close to one then the last positive term  $-\log P(\theta^*)$  is close to zero and the two other terms will dominate, making the whole expression negative, thus, implying convergence. The argument can also be extended to the case where  $D$  is generated by  $\sum_{\xi} Q(D|\xi)Q(\xi)$  instead of  $Q(D|\xi^*)$ , depending again on the weight of  $Q(\xi^*)$  compared to the other weights  $Q(\xi)$ , i.e. how close the other agent is to the Nash policy. The same argument is then repeated for player 2 who uses the prediction model  $Q(a|\xi)$  to model  $P(a|\theta)$ .

Consider as an example the matching pennies game [33], where each player has a penny and must decide whether to secretly turn their penny to heads or tails. Player 1 wins if both pennies show heads or both show tails, whereas player 2 wins if the pennies are unmatched. The payoff matrices of the game are

$$\begin{pmatrix} (+1, -1) & (-1, +1) \\ (-1, +1) & (+1, -1) \end{pmatrix},$$

where the first number in each cell is the payoff for player 1 and the second number is the payoff for player 2, and the columns of the matrix correspond to the choice of heads or tails of player 1, and the rows of the matrix correspond to the choice of heads or tails of player 2. The best responses are then

$$P(a = H|\theta) = \begin{cases} 0 & \text{if } \theta < 1/2 \\ \frac{1}{2} & \text{if } \theta = 1/2 \\ 1 & \text{if } \theta > 1/2 \end{cases}$$

and

$$Q(o = H|\xi) = \begin{cases} 1 & \text{if } \xi < 1/2 \\ \frac{1}{2} & \text{if } \xi = 1/2 \\ 0 & \text{if } \xi > 1/2 \end{cases}$$




Figure 1: Matching pennies game with two Thompson sampling agents. The left figure shows the evolution of the mean of the posterior distributions over  $\theta$  and  $\xi$  for a single run. The right figure shows the average mean reward obtained by each player.

and the predictive models are  $P(o|\theta) = \theta$  and  $Q(a|\xi) = \xi$ . For the priors we assume the uniform distributions  $P(\theta) = 1$  and  $Q(\xi) = 1$ , such that the posteriors are

$$P(\theta; a_1, a_2) = \frac{\theta^{a_1-1}(1-\theta)^{a_2-1}}{B(a_1, a_2)}$$

and

$$Q(\xi; b_1, b_2) = \frac{\xi^{b_1-1}(1-\xi)^{b_2-1}}{B(b_1, b_2)},$$

where  $B(\cdot, \cdot)$  is the beta function, and  $a_1$  and  $a_2$  is the number of heads and tails played by player 2, whereas  $b_1$  and  $b_2$  is the number of heads and tails played by player 1. In this case the only  $(\theta^*, \xi^*)$  pair that is a Nash equilibrium is in the 50 : 50 case, because only then do action and prediction model fit for both agents. In each time step the agents sample  $\theta$  and  $\xi$  respectively from their posteriors  $P(\theta|\cdot)$  and  $Q(\xi|\cdot)$  and act with their best response to this sample. In Figure 4 it can be seen how they co-adapt and converge to the Nash equilibrium.


## 5 Causal induction

Agent (11) can be thought of as a probabilistic superposition of models  $\theta$ , where each model  $\theta$  is characterized by a likelihood model  $P(o_t|\theta, \underline{ao}_{<t}a_t)$  and a policy model  $P(a_t|\theta, \underline{ao}_{<t})$ . In previous applications we assumed that all models  $\theta$  have the same causal structure, i.e. considering multivariate random variables  $a_t$  and  $o_t$ , we assumed that the same variables  $a_t$  are intervened for all  $\theta$  and the same causal model is used to predict the consequences of these interventions on the observational variables  $o_t$ . However, this need not be the case. In principle, different models  $\theta$  could represent different causal structures and

suggest intervention of different variables. Such a setup can be used for causal induction.

Imagine, for example, we are given a device with two light bulbs, one green ( $X$ ) and one red ( $Y$ ), whose states obey a hidden mechanism that correlates them positively. Moreover, the device has switches that allow us controlling the state of either bulb.. We encode the “on” and “off” states of the green light as  $X = x$  and  $X = \neg x$  respectively. Analogously,  $Y = y$  and  $Y = \neg y$  denote the “on” and “off” states of the red light. We are interested in the explanatory power of two competing hypotheses: either “green causes red” ( $\Theta = \theta$ ) or “red causes green” ( $\Theta = \neg\theta$ ).

One of the main methods to deal with problems of causal inference is the framework of causal graphical models [29]. Given a graph that represents a causal structure, we can intervene this graph and ask questions about the probabilities of the variables in the graph. However, in causal induction we would like to discover the causal structure itself, that is we would like to do inference over a multitude of graphs representing different causal structures [38]. If one would like to represent the problem of causal discovery graphically, the main challenge is that the model  $\Theta$  is a random variable that controls the causal structure itself. That is, a tentative graphical representation would be



which cannot be analyzed using the mathematical framework of graphical models alone because the random variable  $\Theta$  operates on a meta-level of the graphical model over  $X$  and  $Y$ . In fact, different causal structures have to be investigated by different graphical models, that is the inference process over different causal structures cannot be represented in one and the same graphical model. However, this difficulty can be overcome by using a probability tree to model the causal structure over the random events [39]. Probability trees can encode alternative causal realizations, and in particular alternative causal hypotheses [40]. All random variables are then of the same type—no distinctions between meta-levels are needed.

## Representation

We can use probability trees to represent the prediction model that the agent has about its environment. An exemplary probability tree for our problem is depicted in Figure 2. In this tree, each (internal) node is interpreted as a causal mechanism; hence a path from the root node to one of the leaves corresponds to a particular sequential realization of causal mechanisms. The logic underlying the structure of this tree is as follows:

1. *Causal precedence:* A node causally precedes its descendants. For instance, the root node corresponding to the sure event  $\Omega$  causally precedes




Figure 2: a) An exemplary probability tree to represent the agent’s prediction model about its environment. The probability of a realization (written under the leaves) is calculated by multiplying the probabilities starting from the root until a leaf is reached. Note that the two hypotheses are statistically indistinguishable. b) The probability tree resulting from (a) after setting  $X = x$ .

all other nodes.

2. *Resolution of variables:* Each node resolves the value of a random variable. For instance, given the node corresponding to  $\Theta = \theta$  and  $X = \neg x$ , either  $Y = y$  will happen with probability  $\Pr(y|\theta, \neg x) = \frac{1}{4}$  or  $Y = \neg y$  with probability  $\Pr(\neg y|\theta, \neg x) = \frac{3}{4}$ .
3. *Heterogeneous order:* The resolution order of random variables can vary across different branches. For instance,  $X$  precedes  $Y$  under  $\Theta = \theta$ , but  $Y$  precedes  $X$  under  $\Theta = \neg\theta$ . This allows modeling different causal hypotheses.

While the probability tree represents our subjective model explaining the order in which the random values are resolved, it does not necessarily correspond to the temporal order in which the events are revealed to us. So for instance, under hypothesis  $\Theta = \theta$ , the value of the variable  $Y$  might be revealed before  $X$ , even though  $X$  causally precedes  $Y$ ; and the hypothesis  $\Theta$ , which precedes both  $X$  and  $Y$ , is never observed.

## Interventions

The importance of interventions to detect causal structure is illustrated in Figure 2, as the observational probabilities are completely symmetric for the two halves of the tree. Suppose we observe that both lights are on. Have we learned anything about their causal dependency? A brief calculation shows that this is not the case because the posterior probabilities are equal to the prior probabilities:

$$\begin{aligned}
 \Pr(\theta|x, y) &= \frac{\Pr(y|\theta, x)\Pr(x|\theta)\Pr(\theta)}{\Pr(y|\theta, x)\Pr(x|\theta)\Pr(\theta) + \Pr(x|\neg\theta, y)\Pr(y|\neg\theta)\Pr(\neg\theta)} \\
 &= \frac{\frac{3}{4} \cdot \frac{1}{2} \cdot \frac{1}{2}}{\frac{3}{4} \cdot \frac{1}{2} \cdot \frac{1}{2} + \frac{3}{4} \cdot \frac{1}{2} \cdot \frac{1}{2}} = \frac{1}{2} = \Pr(\theta).
 \end{aligned}$$



This makes sense intuitively, because by just observing that the two lights are on, it is statistically impossible to tell which one caused the other. The only way to extract causal information is then to intervene, paraphrased as “no causes in, no causes out” [41] or “to find out what happens when you kick the system, you have to kick the system” [42]. Thus, we now repeat our experiment, but this time *we* turn on the green light ( $X = x$ ). We reflect this choice by changing all the mechanisms that resolve the random variable  $X$ , placing all the probability mass on the outcome  $X = x$  (see Figure 2b). Assume that we subsequently observe that the second light is on. Then, the posterior probabilities are

$$\begin{aligned}\Pr(\theta|\hat{x}, y) &= \frac{\Pr(y|\theta, \hat{x})\Pr(\hat{x}|\theta)\Pr(\theta)}{\Pr(y|\theta, \hat{x})\Pr(\hat{x}|\theta)\Pr(\theta) + \Pr(\hat{x}|\neg\theta, y)\Pr(y|\neg\theta)\Pr(\neg\theta)} \\ &= \frac{\frac{3}{4} \cdot 1 \cdot \frac{1}{2}}{\frac{3}{4} \cdot 1 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2} \cdot \frac{1}{2}} = \frac{3}{5},\end{aligned}$$

where  $\hat{x}$  is Pearl’s notation to indicate a causal intervention of  $X$ . Since  $P(\theta) < P(\theta|\hat{x}, y)$ , we have gathered evidence favoring the hypothesis “green causes red”. This was only possible because our intervention introduced a statistical asymmetry among the two hypotheses that did not exist before.

### Thompson Sampling

Naturally, multiple interventions and observations can be executed in consecution. In this case Thompson sampling is used in each time step to decide which policy model to use, which implies the decision which variables to intervene. Then, after the intervention, all variables are revealed simultaneously at every time step of the inference process. The update of the observational probabilities is done the same way as in the one step case, taking into account which variables were intervened. A simulation of the repeated Thompson sampling process for causal induction of our example system is shown in Figure 3.

## 6 Discussion

Equations (11) was first derived in [17] as the optimal solution to the adaptive coding problem given actions and observations as *Bayesian rule for control*. In practice, it is implemented by sampling an environment parameter  $\hat{\theta}_t$  for each time step from the posterior distribution  $B(\theta|\hat{a}_{<t}, o_{<t})$ , and then treating it as if it was the true parameter—that is, issuing the action  $a_t$  from  $B(a_t|\hat{\theta}_t, a_{<t}, o_{<t})$ . This action-sampling method where beliefs are randomly instantiated was first proposed as a heuristic in [1] and is now known as *Thompson sampling*. Equations (11) therefore provides a method for generalized Thompson sampling applicable to adaptive sequential decision-making problems.

The main contribution of this paper is to examine three features of such generalized Thompson sampling. First, we provide an argument showing that Thompson sampling is a natural consequence of a Bayesian treatment of policy uncertainty. Policy uncertainty arises whenever an agent is trying to find an optimal policy, but is unable to do so—for example due to computational




Figure 3: Thompson sampling for causal induction. (Left) Posterior distribution  $P(\theta|\cdot)$  for 10 runs when the true system is given by  $\Theta = \theta$ . (Right) Posterior distribution  $P(\theta|\cdot)$  for 10 runs when the true system is given by  $\Theta = -\theta$ . In both cases the agent is able to identify the causal structure of the environment with high confidence when  $P(\theta|\dots)$  is close to one or zero respectively.

constraints—, even though the agent might have a coarse idea about the optimum, which can be expressed as a probability distribution. The Bayesian treatment of this uncertainty is analogous to Bayesian estimation in the case of pure observation problems. The Bayes-optimal estimator in this case is not point estimate, but a distribution, which forgoes the bias-variance dilemma. Similarly, in the case of actions, the exploration-exploitation trade-off can be circumvented by Thompson sampling from probabilistic policies expressed as Bayesian mixture distributions.

Second, we investigated co-adaptation of two adaptive Thompson sampling agents. We could demonstrate that such agents converge to Nash equilibria, if the parameterized policy set they are choosing from is given by the parameterized best response functions. This approach also generalizes previous models of fictitious play [43, 44] that best-respond to the observed frequency of the opponent’s play rather than best-responding to their randomized beliefs about the opponent. Therefore, adaptive Thompson sampling agents might provide a useful modeling tool for evolutionary game theory [35] and learning in games in general [45].

Third, we could demonstrate that generalized Thompson sampling can also be applied to the problem of causal induction, by designing policy and prediction models with different causal structures. This way generalized Thompson sampling can be used as a general method for causal induction that is Bayesian in nature. It is based on the idea of combining probability trees [40] with interventions [29] for predicting the behavior of a manipulated system with multiple causal hypotheses. Both the interventions and the constraints on the causal hypotheses introduce statistical asymmetries that permit the extraction of causal information. Unlike frameworks that aim to extract causal information from ob-

servational data alone [46, 47, 48], the proposed method is designed for agents that interact with their environment and use these interactions to discover causal relationships.

So far Thompson sampling has been mainly applied to multi-armed bandit problems. Multi-armed bandits can be represented by a parameter  $\theta$  that summarizes the statistical properties of the reward obtained for each lever. Reward distributions range from Bernoulli to Gaussian (with unknown mean and variance), and they can also depend on the particular context or state [9, 11, 8, 10]. In particular, the work of [11] and the work of [8] prove asymptotic convergence of Thompson sampling. The work of [12] presents empirical results that show Thompson sampling is highly competitive, matching or outperforming popular methods such as UCB [25, 26].

Another class of problems, where Thompson sampling has been applied in the past, are Markov decision processes (MDPs). MDPs can be described by parameterizing the dynamics and reward distribution (model-based) [49] or by directly parameterizing the *Q-table* (model-free) [50, 31]. The first approach samples a full description of an MDP, solves it for the optimal policy, and then issues the optimal action. This is repeated in each time step. The second approach avoids the computational overhead of solving for the optimal policy in each time step by directly doing inference on the Q-tables. Actions are chosen by picking the one having the highest Q-value for the current state. The same ideas can also be applied to solve adaptive control problems with linear system equations, quadratic cost functions and Gaussian noise [51].

## 6.1 Optimality

One of the main arguments is that the derivation presented in Section 3.3 shows that generalized Thompson sampling is not just a heuristic method, but that it can be derived under the assumption of policy uncertainty—simply by applying probability theory and causal calculus. This Thompson sampling approach differs from the formulation of adaptive control problems as control problems with known environments that require the maximization of a subjective expected utility criterion—compare Section 2.3. The difference between the two approaches can be highlighted by contrasting the two one-step decision scenarios depicted in Figure 4. The goal is to predict the outcome of a biased coin with payoffs \$1 and \$0 for a right and wrong guess respectively. A rational decision maker places bets (shown inside speech bubbles) such that his subjective expected utility is maximized. These subjective beliefs are delimited within dotted boxes.

The difference between the two becomes clear by inspecting the expected utility in each case: they are

$$\max_{P'} \sum_{\theta} P(\theta) \left\{ \sum_o P'(a|\theta) P(o|\theta, a) U(o) \right\}, \quad (a)$$

$$\text{and } \sum_{\theta} P(\theta) \max_{P'} \left\{ \sum_o P'(a|\theta) P(o|\theta, a) U(o) \right\} \quad (b)$$

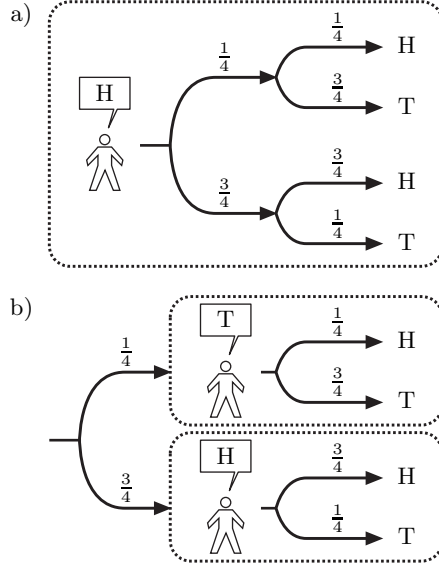


Figure 4: Comparison between (a) a fully rational decision maker and (b) a decision maker with random beliefs.

respectively. Here it is clearly seen that the difference between the two lies in the order in which we apply the expectation (over the environment parameter) and the maximization operator. Both cases can be explained in terms of optimality. However, in (a), decision-maker picks his action taking into account the uncertainty over the bias, while in (b), the decision-maker picks his action only after his beliefs over the coin bias are instantiated—that is, *he is optimal w.r.t. his random beliefs*.

The difference between probabilities that one takes into account when making a decision versus the probabilities that are not (i.e. they are *immeasurable*) has been first proposed by [52]. The classical decision theories of [18] and [19] only consider known probabilities that are reasoned about inside the max-operation. Another example where random beliefs play a crucial role is in *games with incomplete information* [33]. Here, having incomplete information about the other player leads to a infinite hierarchy of meta-reasoning about the other player’s strategy. To avoid this difficulty, Harsanyi introduced *Bayesian games* [53]. In a Bayesian game, incomplete knowledge is modeled by randomly instantiating the player’s types, after which they choose their strategies optimally—thus eliminating the need for recurrent reasoning about the other players’ strategy.

Maintaining and updating Bayesian probabilities is an optimally efficient way to deal with uncertainty—be it with respect to the policy or the environment [31]. Therefore, the central claim is that having random beliefs—as

formalized by generalized Thompson sampling—can be considered optimal *under the constraint* of having policy uncertainty—uncertainty that is inevitable whenever we are unable to compute the optimal policy. Having policy uncertainty effectively weakens the two assumptions of the maximum expected utility principle: the optimal policy can be chosen and refined during interactions, and the computational complexity is lower.

The operational distinction of having policy uncertainty has important algorithmic consequences. When there is policy uncertainty, the belief of the decision-maker is itself a random variable. This means that the very policy is undefined until the random variable is resolved. Hence, the computation of the optimal policy can be delayed and determined dynamically. It is precisely this fact that is (implicitly) exploited in popular reinforcement learning algorithms, and explicitly in the algorithms based on random beliefs. This is in stark contrast to the case when there is no policy uncertainty, where the policy is pre-computed and static.

### Adaptive Coding and the Kullback-Leibler Divergence

Even though the maximization is inside the expectation in case of random belief approaches to decision-making, there is another outer maximization or optimality criterion implicit, analogous to the case of Bayesian inference that is known to optimize Kullback-Leibler divergences. Therefore, it is useful to think about the adaptive control problem as an adaptive coding and inference problem [17, 54]. In terms of the initial problem statement in Section 2, the question then is: How can the designer construct a system  $P$  defined by  $P(o_t|a_{\leq t}, o_{< t})$  and  $P(a_t|a_{< t}, o_{< t})$  such that its behavior is as close as possible to the custom-made system  $B(o_t|\theta, a_{\leq t}, o_{< t})$  and  $B(a_t|\theta, a_{< t}, o_{< t})$  under any realization of  $Q_\theta$ ? Using the Kullback-Leibler divergence as a distance measure, we can formulate a variational problem in  $P$

$$P := \arg \min_{\mathbf{Pr}} \left\{ \limsup_{t \rightarrow \infty} \sum_{\theta} P(\theta) \sum_{\tau=1}^t (D_m^{a_\tau}(\mathbf{Pr}) + D_m^{o_\tau}(\mathbf{Pr})) \right\} \quad (12)$$

with

$$D_\theta^{a_t}(\mathbf{Pr}) = \sum_{a_{< t}, o_{< t}} B(a_{< t}, o_{< t}|\theta) \sum_{a_t} B(a_t|\theta, a_{< t}, o_{< t}) \log \frac{B(a_t|\theta, a_{< t}, o_{< t})}{\mathbf{Pr}(a_t|a_{< t}, o_{< t})}$$

$$D_\theta^{o_t}(\mathbf{Pr}) = \sum_{a_{\leq t}, o_{< t}} B(a_{\leq t}, o_{< t}|\theta) \sum_{o_t} B(o_t|\theta, a_{\leq t}, o_{< t}) \log \frac{B(o_t|\theta, a_{\leq t}, o_{< t})}{\mathbf{Pr}(o_t|a_{\leq t}, o_{< t})}.$$

In the case of observations, this is a well-known variational principle for Bayesian inference, as it describes a predictor that requires, on average, the least amount of extra bits to capture informational surprise stemming from the behavior of the environment. In the case of actions, the same principle can be harnessed to describe resourceful generation of actions in a way that requires random bits with minimum length on average, when trying to match the optimal policy most suitable for the unknown environment [55].

## 6.2 Evolutionary game theory

When dealing with adaptive agents, one of the most intriguing questions is what happens if two adaptive agents are coupled. Classic game theory does not really allow to address that question, as it abstracts away from learning and adaptation processes and focuses on fix point conditions for equilibria. In contrast, evolutionary game theory focuses on the dynamics that can lead to equilibria [35]. One of the most widely studied dynamics in evolutionary game theory are the so-called replicator equations

$$x_{t+1}^i = \frac{x_t^i f_i(x)}{\sum_j x_t^j f_j(x)},$$

where  $x_t^i$  represents the proportion of type  $i$  in a population of individuals at time  $t$ . The vector  $x_t = (x_t^1, \dots, x_t^n)$  represents the population distribution, such that  $\sum_j x_t^j = 1$ . The function  $f_i(x)$  denotes the fitness of type  $i$ , which depends on the population  $x$ . The proportion of individuals of type  $i$  at the next time point  $t + 1$  is determined by the fitness share this type achieves compared to the population total.

Interestingly, there is a formal correspondence between the replicator dynamics and Bayesian inference [36, 37]

$$p(h|d) = \frac{p(h)p(d|h)}{\sum_{h'} p(h')p(d|h')},$$

where  $p(h)$  and  $p(h|d)$  represents the prior and posterior probability mass allotted to hypothesis  $h$ . The likelihood function  $p(d|h)$  works as a fitness landscape. The posterior probability is determined by the likelihood fitness achieved compared to the overall evidence  $P(d) = \sum_{h'} p(h')p(d|h')$ .

In evolutionary game theory the fix points of the replicator dynamics have been studied extensively. In particular, evolutionarily stable strategies [56] have been shown to be a refinement of the common Nash equilibrium, in the sense that such Nash equilibria are stable with respect to perturbing mutant strategies. Since generalized Thompson sampling as described in (11) shares its form with Bayesian inference, the connection to evolutionary replicator equations is immediate. Therefore, we could apply very similar stability arguments in the case of two interacting adaptive agents (11), as previously applied in the case of the replicator dynamics. Generalized Thompson sampling might therefore also provide a useful tool in the future to study convergence of co-adaptation processes within the context of evolutionary game theory.

## 6.3 Causality

To construct the Bayes-causal solution in Section 3.3, we needed to treat actions as interventions. This raises the question about why this distinction was not made for deriving classical SEU solutions. More formally, under what conditions

do the equations in (11) reduce to

$$\begin{aligned} P(a_t|a_{<t}, o_{<t}) &= B(a_t|a_{<t}, o_{<t}) \\ P(o_t|a_{\leq t}, o_{<t}) &= B(o_t|a_{\leq t}, o_{<t}), \end{aligned}$$

that is, with no interventions?

Since,

$$\begin{aligned} B(a_t|a_{<t}, o_{<t}) &= \sum_{\theta} B(a_t|\theta, a_{<t}, o_{<t}) B(\theta|\hat{a}_{<t}, o_{<t}) \\ B(o_t|a_{\leq t}, o_{<t}) &= \sum_{\theta} B(o_t|\theta, a_{\leq t}, o_{<t}) B(\theta|\hat{a}_{\leq t}, o_{<t}), \end{aligned}$$

determining the conditions boils down to analyzing when the equalities

$$\begin{aligned} B(\theta|a_{<t}, o_{<t}) &= B(\theta|\hat{a}_{<t}, o_{<t}) \\ B(\theta|\hat{a}_{\leq t}, o_{<t}) &= B(\theta|a_{\leq t}, o_{<t}) \end{aligned}$$

hold. Replacing both sides yields,

$$\begin{aligned} & \frac{B(\theta) \prod_{k=1}^t B(a_k|\theta, a_{<k}, o_{<k}) B(o_k|\theta, a_{\leq k}, o_{<k})}{\sum_{\theta'} B(\theta') \prod_{k=1}^t B(a_k|\theta', a_{<k}, o_{<k}) B(o_k|\theta', a_{\leq k}, o_{<k})} \\ &= \frac{B(\theta) \prod_{k=1}^t B(o_k|\theta, a_{\leq k}, o_{<k})}{\sum_{\theta'} B(\theta') \prod_{k=1}^t B(o_k|\theta', a_{\leq k}, o_{<k})} \end{aligned} \quad (13)$$

Inspecting (13) we conclude that

$$B(a_k|\theta, a_{<k}, o_{<k}) = \delta_{\bar{a}_k}(a_k),$$

i.e. the actions have to be issued deterministically (but possibly history-dependent) from a unique policy. Intuitively speaking, this is because the operations of intervening and conditioning coincide when the random variables are deterministic.

## 6.4 Open Problems

There are important cases where random belief approaches can fail. Indeed, it is easy to devise experiments where having policy uncertainty converges exponentially slower (or does not converge at all) than knowing the optimal policy.

Consider the following simple example: Environment 1 is a  $k$ -state MDP in which only  $k$  consecutive actions  $A$  reach a state with reward +1. Any interception with a  $B$ -action leads back to the initial state. A second environment which is like the first but actions  $A$  and  $B$  are interchanged. The optimal policy figures out the true environment in  $k$  actions (either  $k$  consecutive  $A$ 's or  $B$ 's). Consider now an agent with random beliefs: The optimal action in environment 1 is  $A$ , in environment 2 is  $B$ . A uniform  $(\frac{1}{2}, \frac{1}{2})$  prior over the two possible environments stays a uniform posterior as long as no reward has been observed. Hence, an agent with random beliefs chooses at each time-step  $A$  and  $B$  with

equal probability. With this policy it takes about  $2^k$  actions to accidentally choose a row of  $A$ 's (or  $B$ 's) of length  $k$ . From then on the agent acts optimally too. Thus, the optimal policy converges in time  $k$ , while the agent with policy uncertainty needs exponentially longer. A simple way to remedy this problem is, of course, to sample random beliefs only every  $k$  time steps. But this problem can be exacerbated in non-stationary environments. Take for instance, an increasing MDP with  $k = \lceil 10\sqrt{t} \rceil$ , in which the optimal policy converges in 100 steps, while an agent with policy uncertainty would not converge at all in most realizations.

Although [17] prove asymptotic convergence for general environments fulfilling a restrictive form of ergodicity condition, this condition needs to be weakened for the convergence proof to be applicable to most real problems. But it is clear that a form of ergodicity is required for an agent with policy uncertainty to be able to learn to act optimally. Intuitively, this means that an agent can only learn if the environment has temporally stable statistical properties. Finally, determining the speed of convergence and the *regret* is currently an open problem.

## 7 Conclusion

In this paper we have argued that policy uncertainty is a natural phenomenon that arises whenever there are not enough computational resources to apply the maximum SEU principle. We have shown that treating this uncertainty in a Bayesian way with actions as random variables that obey causal calculus naturally leads to Thompson sampling and its Bayesian generalization. This generalized Thompson sampling can be straightforwardly applied to evolutionary game theory and to the problem of causal induction. As these random-belief approaches can be derived simply from probability theory and causal calculus we suggest that they should not be considered as mere heuristics but as well-founded principled approaches.

## 8 Acknowledgments

This study was supported by the DFG, Emmy Noether grant BR4164/1-1.

## References

- [1] Thompson WR. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*. 1933;25(3/4):pp. 285–294.
- [2] Myers JL. ‘Probability learning and sequence learning’. In: Estes WK, editor. *Handbook of Learning and Cognitive Processes: Approaches to Human Learning and Motivation*. Hillsdale, NJ: Erlbaum; 1976. p. 171–205.



- [3] Vulkan N. An Economists Perspective on Probability Matching. *Journal of Economic Surveys*. 2000;14(1):101–118.
- [4] Sutton RS, Barto AG. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press; 1998.
- [5] Gittins JC. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society Series B, Methodological*. 1979;41:148–177.
- [6] Granmo OC. A Bayesian Learning Automaton for Solving Two-Armed Bernoulli Bandit Problems. In: *Proceedings of the 2008 Seventh International Conference on Machine Learning and Applications*. ICMLA '08; 2008. p. 23–30.
- [7] Asmuth J, Li L, Littman ML, Nouri A, Wingate D. A Bayesian sampling approach to exploration in reinforcement learning. In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. UAI '09. Arlington, Virginia, United States: AUAI Press; 2009. p. 19–26.
- [8] Granmo OC. Solving two-armed Bernoulli bandit problems using a Bayesian learning automaton. *International Journal of Intelligent Computing and Cybernetics*. 2010;3(2):207–234.
- [9] Graepel T, Quionero Candela J, Borchert T, Herbrich R. Web-Scale Bayesian Click-Through rate Prediction for Sponsored Search Advertising in Microsoft's Bing Search Engine. In: *Proceedings of the Twenty-Seventh International Conference on Machine Learning*; 2010. p. 25–26.
- [10] Scott S. A modern Bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*. 2010;26:639–658.
- [11] May BC, Leslie DS. Simulation studies in optimistic Bayesian sampling in contextual-bandit problems. *Statistics Group, Department of Mathematics, University of Bristol*; 2011. 11:02.
- [12] Chapelle O, Li L. An empirical evaluation of Thompson sampling. *NIPS*. 2011;.
- [13] Agrawal S, Goyal N. Analysis of Thompson sampling for the multi-armed bandit problem. *NIPS Workshop paper*. 2011;.
- [14] Granmo OC, Glimsdal S. Accelerated Bayesian learning for decentralized two-armed bandit based decision making with applications to the Goore Game. *Applied Intelligence*. 2012;.
- [15] May BC, Korda N, Lee A, Leslie DS. Optimistic Bayesian Sampling in Contextual-Bandit Problems. *J Mach Learn Res*. 2012 Jun;98888:2069–2106.
- [16] Russo D, Roy BV. Learning to Optimize Via Posterior Sampling. *arXiv*. 2013;abs/1301.2609.

- [17] Ortega PA, Braun DA. A minimum relative entropy principle for learning and acting. *Journal of Artificial Intelligence Research*. 2010;38:475–511.
- [18] Von Neumann J, Morgenstern O. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press; 1944.
- [19] Savage LJ. *The Foundations of Statistics*. New York: John Wiley and Sons; 1954.
- [20] Bellman RE. *Dynamic Programming*. Princeton, NJ: Princeton University Press; 1957.
- [21] Simon H. A Behavioral Model of Rational Choice. In: *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. Wiley; 1957. .
- [22] Hutter M. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Berlin: Springer; 2004.
- [23] Duff MO. *Optimal learning: computational procedures for bayes-adaptive markov decision processes*; 2002. Director-Andrew Barto.
- [24] Pearl J. A Causal Calculus for Statistical Research. In: Fisher D, Lenz HJ, editors. *Learning from Data: AI and Statistics V*. Springer-Verlag; 1996. .
- [25] Lai TL, Robbins H. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*. 1995;6:4–22.
- [26] Auer P, Cesa-Bianchi N, Fisher P. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*. 2002;47:235–256.
- [27] Geman S, Bienenstock E, Doursat R. Neural Networks and the Bias/Variance Dilemma. *Neural Computation*. 1992;4:1–58.
- [28] Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian Data Analysis*. Chapman and Hall/CRC; 2003.
- [29] Pearl J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK; 2009.
- [30] Spirtes GC P, Scheines R. *Causation, Prediction, and Search*. 2nd ed. MIT Press; 2000.
- [31] Ortega PA, Braun DA. A Bayesian rule for adaptive control based on causal interventions. In: *Proceedings of the third conference on general artificial intelligence*; 2010. .
- [32] Gibbons R. *A Primer in Game Theory*. Financial Times Prentice Hall; 1992.
- [33] Osborne MJ, Rubinstein A. *A Course in Game Theory*. MIT Press; 1999.

- [34] Nash JF. Equilibrium Points in n-Person Games. Proceedings of the National Academy of Sciences of the United States of America. 1950 jan;36(1):48–49.
- [35] Weibull JW. Evolutionary Game Theory. Cambridge, MA: MIT Press; 1995.
- [36] Shahlizi CR. Dynamics of bayesian updating with dependent data and misspecified models. Electronic Journal of Statistics. 2009;3:1039–1074.
- [37] Harper M. The Replicator Equation as an Inference Dynamic. arXiv. 2009;abs/0911.1763.
- [38] Heckerman D, Meek C, Cooper G. A Bayesian Approach to Causal Discovery. Microsoft Research Technical Report MSR-TR-97-05; 1997.
- [39] Ortega PA. Bayesian Causal Induction. NIPS Workshop on Philosophy and Machine Learning. 2011;.
- [40] Shafer G. The Art of Causal Conjecture. MIT Press; 1996.
- [41] Cartwright N. Nature’s Capacities and Their Measurements. Oxford University Press; 1994.
- [42] Box GEP. Use and Abuse of Regression. Technometrics. 1966;8(4):625–629.
- [43] Brown GW. Iterative solutions of games by fictitious play. In: Koopmans TC, editor. Activity Analysis of Production and Allocation. Wiley, NY; 1951. .
- [44] Fudenberg D, Levine D. The Theory of Learning in Games. MIT Press, Cambridge, MA; 1998.
- [45] Fudenberg D, Kreps D. Learning mixed equilibria. Games and Economic Behavior. 1993;5:320–367.
- [46] Shimizu S, Hoyer PO, Hyvärinen A, Kerminen A. A Linear Non-Gaussian Acyclic Model for Causal Discovery. J Mach Learn Res. 2006 Dec;7:2003–2030.
- [47] Griffiths TL, Tenenbaum JB. Theory-based causal induction. Psychological Review. 2009;116:661–716.
- [48] Janzing D, Schölkopf B. Causal inference using the algorithmic Markov condition. IEEE Trans Inf Theor. 2010 Oct;56(10):5168–5194.
- [49] Strens M. A Bayesian Framework for Reinforcement Learning. In: Proceedings of the Seventeenth International Conference on Machine Learning; 2000. .

- [50] Dearden R, Friedman N, Russell S. Bayesian Q-learning. In: AAAI '98/IAAI '98: Proceedings of the fifteenth national/tenth conference on Artificial intelligence/Innovative applications of artificial intelligence. Menlo Park, CA, US: American Association for Artificial Intelligence; 1998. p. 761–768.
- [51] Braun DA, Ortega PA. A minimum relative entropy principle for adaptive control in linear quadratic regulators. In: The 7th conference on informatics in control, automation and robotics. vol. 3; 2010. p. 103–108.
- [52] Knight FH. Risk, Uncertainty, and Profit. Boston: Houghton Mifflin; 1921.
- [53] Harsanyi J. Games with Incomplete Information Played by "Bayesian" Players. Management Science. 1967;14(3):159–182.
- [54] Ortega PA, Braun DA. Adaptive Coding of Actions and Observations. NIPS Workshop on Information in Perception and Action. 2012;.
- [55] MacKay DJC. Information Theory, Inference, and Learning Algorithms. Cambridge University Press; 2003.
- [56] Smith JM. Evolution and the Theory of Games. Cambridge, UK: Cambridge University Press; 1982.